

# CARE: Confidence-Rich Autonomous Robot Exploration Using Bayesian Kernel Inference and Optimization

Yang Xu <sup>1</sup>, Student Member, IEEE, Ronghao Zheng <sup>1</sup>, Member, IEEE, Senlin Zhang <sup>1</sup>, Member, IEEE, Meiqin Liu <sup>1</sup>, Senior Member, IEEE, and Shoudong Huang <sup>1</sup>, Senior Member, IEEE

**Abstract**—In this letter, we consider improving the efficiency of information-based autonomous robot exploration in unknown and complex environments. We first utilize Gaussian process (GP) regression to learn a surrogate model to infer the confidence-rich mutual information (CRMI) of querying control actions, then adopt an objective function consisting of predicted CRMI values and prediction uncertainties to conduct Bayesian optimization (BO), i.e., GP-based BO (GPBO). The trade-off between the best action with the highest CRMI value (exploitation) and the action with high prediction variance (exploration) can be realized. To further improve the efficiency of GPBO, we propose a novel lightweight information gain inference method based on Bayesian kernel inference and optimization (BKIO), achieving an approximate logarithmic complexity without the need for training. BKIO can also infer the CRMI and generate the best action using BO with bounded cumulative regret, which ensures its comparable accuracy to GPBO with much higher efficiency. Extensive numerical and real-world experiments show the desired efficiency of our proposed methods without losing exploration performance in different unstructured, cluttered environments.

**Index Terms**—Probabilistic inference, range sensing, reactive and sensor-based planning, view planning for SLAM.

## I. INTRODUCTION

ROBOT exploration gains its prevalence recently in priori unknown environments such as subterranean, marine, and

Manuscript received 17 May 2023; accepted 1 September 2023. Date of publication 7 September 2023; date of current version 14 September 2023. This letter was recommended for publication by Associate Editor M. Popovic and Editor J. Civera upon evaluation of the reviewers' comments. This work was supported in part by the NSFC-Zhejiang Joint Fund for the Integration of Industrialization and Informatization under Grant U1909206 and in part by the National Natural Science Foundation of China under Grant 62173294. (Corresponding author: Ronghao Zheng.)

Yang Xu, Ronghao Zheng, and Senlin Zhang are with the College of Electrical Engineering, Zhejiang University, Hangzhou 310027, China, and also with the State Key Laboratory of Industrial Control Technology, Zhejiang University, Hangzhou 310027, China (e-mail: xuyang94@zju.edu.cn; rzheng@zju.edu.cn; slzhang@zju.edu.cn).

Meiqin Liu is with the Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University, Xi'an 710049, China, and also with the State Key Laboratory of Industrial Control Technology, Zhejiang University, Hangzhou 310027, China (e-mail: liumeiqin@zju.edu.cn).

Shoudong Huang is with the Robotics Institute, Faculty of Engineering and Information Technology, University of Technology Sydney, Ultimo, NSW 2007, Australia (e-mail: shoudong.huang@uts.edu.au).

We also provide our open-source implementation code at <https://github.com/Shepherd-Gregory/BKIO-Exploration>.

Digital Object Identifier 10.1109/LRA.2023.3313054




Fig. 1. CRMI-based autonomous robot exploration using a range-sensing Fetch robot with limited FOV in a cluttered office environment. **Left:** Fetch robot. **Middle:** Environmental CRMI surface after exploration. It shows the resulting exploration trajectory (white) in one experiment, where the robot's overall exploration sequence to maximize the CRMI is A-B-G-C-D-E-F-D-A-B-G-A. **Right:** Corresponding real-world views at points D and F. The color bar is the MI scale [0, 1] bit. Note that this informative trajectory is planned *online* by the CRMI-driven planner. The CRMI-driven robot will explore the unexplored spaces (*dark red*), then obscured spaces (*light red*), and be repulsive to walls/obstacles (*deep blue*). The gray lines represent glass rooms, walls, or building boundaries the robot can not enter.

planetary tasks [1], [2]. Among the literature, state-of-the-art exploration methods prefer to use information-theoretic metrics in each iteration, such as Shannon mutual information (MI) [3] and its derivatives [4], [5], [6], [7], to evaluate the information gain brought by candidate control actions accurately. For instance, in [3], Julian et al. defined the Shannon MI between new observation and the occupancy grid map (OGM) at candidate poses. Charrow et al. [4] adopted the Cauchy-Schwarz quadratic MI (CSQMI) to improve the time efficiency of Shannon MI in 2D and 3D exploration for computationally limited platforms. These OGM-based MI metrics are also called occupancy grid MI (OGMI). Our previous work [6] further provided a more accurate information measure, confidence-rich mutual information (CRMI), over the dense confidence-rich map (CRM) [8] for robot exploration.

Driven by these information metrics, the robot will choose and execute the most informative action, thus the exploration problem becomes a sequential optimal decision-making one naturally. According to [3], the good properties of MI also can ensure the accuracy and the eventual completeness of the exploration. A real-world exploration example is in Fig. 1.

Intuitively, the way to tackle an autonomous exploration problem is to use a greedy strategy and add more candidate actions, including sampled nodes [9], [10], available viewpoints [11], or special motion primitives [12], in the discrete action space.

However, the computational cost of the MI evaluation of all candidate actions will become expensive since the forward simulation in the evaluation requires extensive raycasting, map updates, and MI calculation. Notably, these consequences will be more distinct in 3D environments because the increased dimension needs to evaluate much more samples.

### A. Related Work

In the context of robot exploration, deep learning has been introduced to realize predicting optimal sensing actions more efficiently. Bai et al. [13] trained the deep neural network with plenty of randomly generated 2D gray maps to generate suggested action and ensure inferring in constant time. Graph neural networks have also been combined with reinforcement learning methods to learn the best action from an exploration graph, rather than metric maps or visual images [14], [15]. Nevertheless, the neural network-based robot exploration methods require numerous offline training samples beforehand and are limited to the adaptability and generalization capability in different environments.

More recently, statistical learning techniques also provided a powerful tool to find the global optimum approximately by training predictive models using minor parts of actions in continuous action spaces without evaluating the objective function expensively, which also has better interpretability in black-box inference [16], [17], [18]. In [19], Bai et al. used the Gaussian process (GP) to model the relationship between control actions and the explicitly evaluated OGMI for the robot exploring priori unknown areas. In [20], they introduced Bayesian optimization (BO) into the OGMI-based robot exploration to optimize the GP prediction in multiple iterations.

However, the main bottleneck of the GP-based methods is that the number of training actions  $N$  will affect the resulting prediction accuracy directly, as well as the computational cost. That implies one needs to pay expensive computations to achieve higher exploration performance. Typically, updating and querying the GP models (the engine behind BO) have an overall  $\mathcal{O}(N^3)$  time complexity. This compromises the inference efficiency and real-time performance of robot exploration tasks inevitably, especially in larger and 3D scenes.

Notably, the Bayesian kernel inference (BKI) technique proposed in [21] gives us a chance to perform efficient exact inference on a simplified model, rather than approximating inference on an exact generative model (e.g. GP) expensively. BKI extends local kernel estimation to Bayesian inference for exponential likelihood functions, enabling only  $\mathcal{O}(\log N_q)$  ( $N_q$ : the number of querying samples) run time for inference. These significant merits enhance BKI's application in robotics, including sensor uncertainty estimation [22], high-speed navigation [23], as well as environment mapping using sparse sensor measurements such as terrain traversability mapping [24], 3D occupancy mapping [25], semantic mapping [26].

### B. Motivations and Contributions

Motivated by the recent BKI technique [21] and its robotic applications [24], [25], [26], in this letter, we aim to utilize it to accelerate the computationally expensive CRMI evaluation of control actions considerably without losing CRMI prediction accuracy and exploration performance in complex scenes and

use BO to optimize the decision-making process. Our main contributions are three-fold:

- 1) We first present a regular CRMI inference method based on GP and BO (GPBO) using minor explicitly evaluated samples to predict CRMI values and uncertainties of candidate actions, and use an information-theoretic objective function to realize the trade-off of the exploration-exploitation dilemma in MI prediction, rather than the naive-greedy (NG) method evaluating all candidate actions exhaustively and expensively;
- 2) We further propose a novel Bayesian kernel inference and optimization method (BKIO) for CRMI inference without training a surrogate evaluation model. Compared with the GPBO method, BKIO can perform much more efficient MI values and uncertainties prediction in approximate logarithm time with comparable accuracy and exploration performance;
- 3) We conduct extensive numerical simulations in synthetic and dataset maps, and real-world experiments in a cluttered office scene to validate our proposed methods. We also release an open-source implementation of our proposed methods.<sup>1</sup>

The letter's organization is as follows. We formulate the GPBO-based CRMI inference problem in Section II and present BKIO for CRMI in Section III. Simulation results using synthetic data and the real-world experimental results, as well as the discussion, are given in Section IV, followed by conclusions in Section V.

## II. BO-BASED CRMI-DRIVEN ROBOT EXPLORATION

In this letter, for simplicity of discussion, we mainly consider the CRMI-based information-theoretic exploration using a mobile robot equipped with a beam-based range sensor of a limited field of view (FOV) in 2D environments. We aim to efficiently evaluate the exact CRMI of unknown robot configurations sampled in the action space.

### A. Information-Theoretic Exploration Using CRMI

Here we begin with some related definitions and notions. Generally, the robot generates a set of candidate actions  $\mathcal{X}_{\text{action}}$  in the robot's feasible configuration space  $\mathcal{X} \subseteq SE(2)$ . For a candidate configuration  $x_i = [p_i^x, p_i^y, \psi_i] \in \mathcal{X}_{\text{action}}$ ,  $p_i^x$  and  $p_i^y$  denote the robot's position on the map, and  $\psi_i$  denotes the heading angle of the robot. We assume this configuration space has been discretized by a fixed resolution over the 2D static grid map comprised of  $N_m$  grid cells. The continuous occupancy values over this map are  $m = \{m^{[1]}, \dots, m^{[\xi]}, \dots, m^{[N_m]}\}$  ( $m^{[\xi]} \in [0, 1]$ ), i.e., the occupancy level over the independent grid cells, which can be updated and queried by the CRM method.

In a typical CRM, the map belief  $b_t^{m^{[\xi]}} := p(m^{[\xi]} | z_{1:t}, x_{1:t})$  of a cell  $m^{[\xi]}$  at time  $t$  is a non-parametric continuous occupancy distribution conditioned upon the sensor observations  $z_{1:t}$  and robot poses  $x_{1:t}$ . The expected map occupancy over a cell  $m^{[\xi]}$  can be defined as the mathematical expectation of the map belief  $b^{m^{[\xi]}}: \hat{m}^{[\xi]} := \mathbb{E}[m^{[\xi]}] = \int_0^1 m^{[\xi]} b^{m^{[\xi]}} dm^{[\xi]}$ , and the confidence of map occupancy can be defined as the variance of map belief distribution. The occupancy distribution of an unobserved map

<sup>1</sup>[Online]. Available: <https://github.com/Shepherd-Gregory/BKIO-Exploration>

cell  $m^{[\xi]}$  is assumed to be uniform, i.e.,  $\hat{m}^{[\xi]} = 0.5$ . This unobserved cell contributes an entropy of 1 b, and a totally known cell (i.e.,  $\hat{m}^{[\xi]} = 0$  or 1) contains 0-bit entropy according to the information theory.

From the view of information theory, the expected CRMI of a candidate configuration  $x_i$  can be evaluated by the current map entropy and conditional entropy, i.e., the *virtual* measurement of a scan  $z'$  raycasting on currently built map  $m$  at pose  $x_i$ , then updating a virtual map  $m'$  and decomposing the CRMI of a scan into each cell on each independent beam. Thus the CRMI can be accumulated over all cells on the map approximately:

$$I(x_i) = H(m) - H(m'|x_i, z') \approx \sum_{j=1}^{N_z} \sum_{k \in \mathcal{Z}_{[j]}} I(m'_{[k]}; z'_{[j]}) \quad (1)$$

where  $N_z$  is the number of beams of the virtual scan  $z'$ ,  $\mathcal{Z}_{[j]}$  is the index set of map cells intersected with the  $j$ th beam  $z'_{[j]}$  of scan  $z'$ ,  $m'_{[k]}$  is the  $k$ th cell of the virtual map  $m'$ . For more details about CRMI calculation, please see [6] or [27].

The aim of information-theoretic robot exploration is to select the best action  $x_{\text{best}}$  maximizing the expected CRMI:

$$x_{\text{best}} = \underset{x_i \in \mathcal{X}_{\text{action}}}{\operatorname{argmax}} I(x_i). \quad (2)$$

**Proposition 1 ([6]):** The CRMI calculation process involves numerous numerical integration and accumulation operations. Hence, the overall computational complexity of evaluating a configuration is about  $O(N_z N_c^2)$  for updating the map and  $O(N_z N_c^2 \lambda_z^{-1} \lambda_m^{-1})$  for the MI calculation, where  $\lambda_z$  and  $\lambda_m$  are the numerical integration resolutions,  $N_c$  is the maximum size of  $\mathcal{Z}_{[j]}$ .

The explicit CRMI evaluation of all sampled actions imposes notable computational costs for robot exploration, especially in unstructured, cluttered, and 3D scenes, since the robot needs to assess more actions to keep safety and exploration performance in these cases.

### B. BO-Based CRMI Inference for Robot Exploration

To alleviate the computational burden of evaluating numerous actions, here we first introduce GP and BO to infer the CRMI with a training set of explicitly evaluated actions.

Consider a supervised learning-based inference problem on predictive stochastic models  $p(\mathbf{y}|\mathbf{x})$  given a sequence of  $N$  explicitly evaluated samples  $\mathcal{D} = \{(\mathbf{x} = \{x_i\}, \mathbf{y} = \{y_i\})\}_{i=1}^N$ , where  $\mathbf{x}$  and  $\mathbf{y}$  represent the set of evaluated configurations and the resulting CRMI values  $I(m; \mathbf{x})$ , respectively. The main objective is to infer the posterior distribution  $p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D})$  to evaluate CRMI values  $\mathbf{y}^* \in \mathbb{R}^{N_q}$  of the  $N_q$  querying sample inputs  $\mathbf{x}^* \in \mathcal{X}_{\text{action}}$ .

Here we assume  $p(\mathbf{y}|\mathbf{x})$  follows a GP, according to [28], we can estimate the posterior mean  $\bar{\mathbf{y}}^*$  and covariance  $\operatorname{cov}(\mathbf{y}^*)$  of the output of the querying set of candidate actions  $\mathbf{x}^*$ , i.e.:

$$\begin{aligned} \bar{\mathbf{y}}^* &= k(\mathbf{x}^*, \mathbf{x})[k(\mathbf{x}, \mathbf{x}) + \sigma_N^2 \mathbf{I}]^{-1} \mathbf{y}, \\ \operatorname{cov}(\mathbf{y}^*) &= k(\mathbf{x}^*, \mathbf{x}^*) - k(\mathbf{x}^*, \mathbf{x})[k(\mathbf{x}, \mathbf{x}) + \sigma_N^2 \mathbf{I}]^{-1} k(\mathbf{x}, \mathbf{x}^*), \end{aligned} \quad (3)$$

where  $\sigma_N^2$  is the Gaussian noise variance vector associated with input  $\mathbf{x}$ ,  $k(\cdot, \cdot)$  is a kernel function measuring the proximity between two elements in feature space.

Among the available kernel functions, we prefer Mat é rn kernel for its capability of handling sudden transitions of terrain [28], [29] since the potential obstacles and unknown structures in application scenes that have never been seen before will vary the MI values greatly. In practice, we choose a Mat é rn 5/2 kernel with the form as:

$$k(\mathbf{x}^*, \mathbf{x}) = \left(1 + \frac{\sqrt{5}r}{\ell} + \frac{5r^2}{3\ell^2}\right) \exp\left(-\frac{\sqrt{5}r}{\ell}\right), r = \|\mathbf{x}^* - \mathbf{x}\|, \quad (4)$$

where  $\ell$  is a positive characteristic length scale.

At each exploration step, the information-driven robot expects to choose the actions with high predicted CRMI values (*exploitation*) to maximize the information gain locally (i.e. Eq. (2)). Still, the results directly predicted by GP tend to recommend sub-optimal action because of the insufficient evaluated samples and the resulting inaccurate predictive model. Hence, the actions with high predicted uncertainty can also not be neglected in the inference process (*exploration*).

To handle this, we incorporate the prediction uncertainty of CRMI values to the objective function in (2) to realize a local trade-off between exploration and exploitation. Referring to the GP-UCB in [30], here we can get the suggested action maximizing the information objective function  $f^I$  based on (2) and (3) after multiple epochs of optimization:

$$x_s = \underset{\mathbf{x} \in \mathcal{X}_{\text{action}}}{\operatorname{argmax}} f^I, f^I = \bar{\mathbf{y}} + \alpha^{1/2} \sigma_{\mathbf{y}}, \quad (5)$$

where  $\alpha$  is the trade-off factor,  $\sigma_{\mathbf{y}}$  is the standard deviation of CRMI prediction.

For such a sequential multi-armed bandit problem [31], one can use the expected regret to ensure that the multi-epoch optimization is reasonable and practical, i.e. the regret of the upper confidence bound (UCB) function should be bounded over rounds.

**Lemma 1 ([30]):** For finite  $|\mathcal{D}| < \infty$ , let  $\delta \in (0, 1)$  and  $\alpha = 2 \log(|\mathcal{D}|T^2\pi^2/6\delta)$ , the UCB with  $\alpha$  for a sample of a GP with mean function zero and covariance function  $k(\mathbf{x}^*, \mathbf{x})$ , can obtain a regret bound of  $\mathcal{O}^*(\sqrt{N_{\text{epoch}}\gamma_T \log |\mathcal{D}|})$  with high probability, where  $N_{\text{epoch}}$  is the optimization rounds,  $T \in [1, N_{\text{epoch}}]$  is the current optimization round, and  $\gamma_T$  is the max information gain after  $N_{\text{epoch}}$  rounds.

**Remark 1:** According to Lemma 1, the GP-based BO using the UCB of (5) and the Mat é rn 5/2 kernel can achieve the desired trade-off of exploration and exploitation after  $N_{\text{epoch}}$  rounds optimization.

## III. BAYESIAN KERNEL INFERENCE AND OPTIMIZATION FOR CRMI-BASED ROBOT EXPLORATION

In this section, we present a novel and highly efficient CRMI inference method based on the BKI and BO for autonomous robot exploration.

### A. Bayesian Generalized Kernel Inference

Here we continue to use the dataset  $\mathcal{D}$  and the querying set  $\mathbf{x}^*$  in Section II-B, and assume the CRMI output  $\mathbf{y}^*$  and the actions input  $\mathbf{x}^*$  follows the relationship of  $p(\mathbf{y}|\mathbf{x})$  (not a GP model). According to the BKI theory and applications [21], [24], [32], the Bayesian generalized kernel inference problem can be formulated into three steps:

1) *Target parameters inference*: First, this problem can be solved by associating latent parameters  $\theta = \{\theta_i\}_{i=1}^N \in \Theta$  with input  $\mathbf{x}$  in the latent space  $\Theta$ , where the likelihood  $p(\mathbf{y}|\theta)$  is known. Thus the inference on  $\mathbf{y}^*$  can be formulated as an inference on target parameters  $\theta^* \in \Theta$  related to  $\mathbf{x}^*$ :

$$p(\mathbf{y}^*|\mathbf{x}^*, \mathcal{D}) = \int_{\Theta} p(\mathbf{y}^*|\theta^*)p(\theta^*|\mathbf{x}^*, \mathcal{D})d\theta^*, \quad (6)$$

where the posterior distribution of the latent variables can be characterized using Bayes' rule:  $p(\theta^*|\mathbf{x}^*, \mathcal{D}) \propto \int_{\Theta} \prod_{i=1}^N p(y_i|\theta_i)p(\theta_{1:N}, \theta^*|\mathbf{x}_{1:N}, \mathbf{x}^*)d\theta_{1:N}$ .

2) *Latent variables marginalization*: By strongly assuming latent parameters  $\theta_{1:N}$  are conditionally independent given the target parameters  $\theta^*$ :  $p(\theta_{1:N}, \theta^*|\mathbf{x}_{1:N}, \mathbf{x}^*) = \prod_{i=1}^N p(\theta_i|\theta^*, x_i, \mathbf{x}^*)p(\theta^*|\mathbf{x}^*)$ , one can marginalize the latent variables  $\theta_{1:N}$  and then obtain  $p(\theta^*|\mathbf{x}^*, \mathcal{D}) \propto \prod_{i=1}^N p(y_i|\theta^*, x_i, \mathbf{x}^*)p(\theta^*|\mathbf{x}^*)$ .

3) *Kernel approximation*: BKI further defines a distribution that has a special smoothness constraint and bounded Kullback-Leibler divergence  $D_{KL}(g||f)$  between the extended likelihood  $p(y_i|\theta^*, x_i, \mathbf{x}^*)$  represented by  $g$  and the likelihood  $p(y_i|\theta_i)$  represented by  $f$ , i.e., the maximum entropy distribution  $g$  satisfying  $D_{KL}(g||f) \leq \rho(\mathbf{x}^*, \mathbf{x})$  has the form  $g(\mathbf{y}) \propto f(\mathbf{y})^{k(\mathbf{x}^*, \mathbf{x})}$ , where  $\rho(\cdot, \cdot): \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}^+$  is a smoothness bound and  $k(\cdot, \cdot): \mathcal{X} \times \mathcal{X} \rightarrow [0, 1]$  is the kernel function which can be uniquely determined by  $\rho$ . Substituting it into (6) yields:

$$p(\theta^*|\mathbf{x}^*, \mathcal{D}) \propto \prod_{i=1}^N p(y_i|\theta^*)^{k(\mathbf{x}^*, \mathbf{x})}p(\theta^*|\mathbf{x}^*). \quad (7)$$

Thus, the posterior distribution can be exactly inferred using the likelihood from the exponential family (e.g., Gaussian distribution) and assuming the corresponding conjugate prior.

### B. BKI Optimization for Autonomous Robot Exploration

Here we present the BKIO method for information-driven robot exploration, using BO to optimize BKI results in multiple epochs.

1) *Bayesian Kernel CRMI Inference*: According to Section III-A, we assume the underlying likelihood  $p(\mathbf{y}|\theta)$  follows a Gaussian distribution with an unknown mean vector  $\theta$  and a fixed, known covariance matrix function  $\Sigma(\mathbf{x})$ :

$$\mathbf{y} \sim \mathcal{N}(\theta, \Sigma(\mathbf{x})), \Sigma(\mathbf{x}) = \text{diag}(\sigma_\theta^2) \in \mathbb{R}^{N \times N}, \quad (8)$$

where  $\sigma_\theta^2$  is the known variance associated with  $\theta$ . Thus the conjugate prior of  $p(\mathbf{y}|\theta)$  can also be described by a Gaussian distribution using the hyperparameter  $\zeta$  and samples input  $\mathbf{x}$ :

$$p(\theta|\mathbf{x}) = \mathcal{N}(\theta_0, \Sigma(\mathbf{x})/\zeta), \quad (9)$$

where  $\theta_0$  and  $\zeta$  are the initial belief of the mean and the uncertainty of that belief, respectively.  $\theta_0$  is a vector consisting of  $N$  known equal constants  $\theta_0$ , and  $\zeta = 0$  means no confidence and  $\zeta \rightarrow \infty$  indicates full prior knowledge. Here we assume  $\zeta$  is a quite small positive constant since we do not have much prior information about the belief when exploring unknown areas. Therefore, since the querying samples set  $\mathcal{D}^* = \{(\mathbf{x}^* = \{x_i^*\}, \mathbf{y}^* = \{y_i^*\})\}_{i=1}^{N_q}$  also follows the distributions of (8) and (9), given the set of explicitly evaluated samples  $\mathcal{D}$ , substituting

(9) and (8) into (7) yields:

$$p(\theta^*|\mathbf{x}^*, \mathcal{D}) \propto \prod_{i=1}^N \exp\left(-\frac{(y_i - \theta_i)^2}{2\sigma_\theta^2}k(\mathbf{x}^*, \mathbf{1}_{N_q} \times x_i)\right) \cdot \exp\left(-\frac{1}{2} \frac{(\theta_i - \theta_0)^2}{\sigma_\theta^2} \zeta\right), \quad (10)$$

thus the predicted mean and covariance of the CRMI posterior can be derived as follows:

$$\bar{\mathbf{y}}^* = \mathbb{E}[\mathbf{y}^*|\mathbf{x}^*, \mathcal{D}] = \mathbb{E}[\theta^*|\mathbf{x}^*, \mathcal{D}] = \frac{\bar{\mathbf{y}} + \zeta\theta_0}{\zeta + \bar{k}} \simeq \frac{\bar{\mathbf{y}}}{\bar{k}},$$

$$\text{cov}(\mathbf{y}^*) = \mathbb{V}[\theta^*|\mathbf{x}^*, \mathcal{D}] = \frac{\Sigma(\mathbf{x}^*)}{\zeta + \bar{k}} \simeq \frac{\Sigma(\mathbf{x}^*)}{\bar{k}}, \quad (11)$$

where  $\bar{\mathbf{y}}$  and  $\bar{k}$  can be computed by the kernel function:

$$\bar{k} = \sum_{i=1}^N k(\mathbf{x}^*, \mathbf{x}), \bar{\mathbf{y}} = \sum_{i=1}^N k(\mathbf{x}^*, \mathbf{x})y_i. \quad (12)$$

Therefore, given a set  $\mathcal{D}$  of explicitly evaluated samples as the input dataset, we can easily compute the CRMI and the corresponding prediction uncertainty for the querying configurations  $\mathbf{x}^*$  by using (11) and (12).

2) *BKI Optimization*: To optimize the BKI inference results of CRMI, we apply Bayesian optimization and use the objective function (5) to ensure the trade-off between exploration and exploitation at each step.

*Remark 2*: According to [30], as a special finite case of GP, the multivariate Gaussian posterior over  $\mathbf{y}^*$  inferred by BKIO (c.f. Eq. (11)) can also ensure the Bayesian optimization regret is bounded using the Matérn 5/2 kernel function and the UCB function of  $f^I$  in (5), which improves the optimality of final suggested action after multiple epochs optimization in each exploration step.

The autonomous exploration framework based on the BKIO method is given in Algorithm 1, where the function ‘Generations’ (Line 5) generates  $N + N_q$  actions uniformly in the robot’s current sensing range, ‘ComputeCRMI’ (Line 12) refers to (1), ‘BKIOptimization’ (Line 16) is in the Algorithm 2. In each iteration, the robot first generates feasible actions  $x_{\text{cand}}$  in its sensing range, e.g., using frontier-based or sampling-based methods, then computes the CRMI values of the actions  $\mathbf{x}$  randomly sampled from  $x_{\text{cand}}$  (Lines 5–14). The robot uses Algorithm 2 to obtain the sets of best actions and the associated MI values and chooses the action with the maximum MI if it’s higher than the threshold; otherwise, the robot will go back to the last action to end the expansion in this branch (Lines 16–23). Once the best action is confirmed, the robot will use the classic A\* algorithm to plan the local collision-free path from the current pose to the best action, and then update the global map using the measurements and poses along the planned local path (Lines 25–26). The exploration will end when the iteration counts meet  $N_{\text{loop}}$  or the action history set is empty.

In Algorithm 2, for each optimization epoch, we first compute the kernel function matrix and use it to compute the CRMI values with the covariance of the querying set (Lines 4–7). Then we can obtain the objective function values of all querying actions and the action with the highest objective function value (Lines 8–9). We can add it into the best action set if it is already in  $\mathcal{D}$ ; otherwise, we will evaluate it explicitly using (1) and add it into  $\mathcal{D}$  to enlarge the reference dataset and enhance the prediction accuracy (Lines 10–18).

---

**Algorithm 1: BKIO-Exploration( ).**


---

**Require:** Occupancy map at  $k$ th time step  $m_k$ , previous robot poses  $x_{\text{hist}} = x_{0:k-1}$  and current pose  $x_k$ , the number of explicitly evaluated samples  $N$ , information threshold  $I_{\text{th}}$ , the number of querying samples  $N_q$ , loop limit  $N_{\text{loop}}$

- 1: iter = 0
- 2: **while**  $x_{\text{hist}} \neq \emptyset$  OR iter <  $N_{\text{loop}}$  **do**
- 3: iter = iter + 1;
- 4: // Generate  $N + N_q$  actions uniformly near  $x_k$
- 5:  $\mathbf{x}_{\text{cand}} \leftarrow \text{GenActions}(x_k, m_k)$ ;
- 6: // Sample  $N$  actions randomly from  $\mathbf{x}_{\text{cand}}$
- 7:  $\mathbf{x} \leftarrow \text{RandSample}(\mathbf{x}_{\text{cand}}, N)$ ;
- 8:  $\mathbf{x}^* \leftarrow \mathbf{x}_{\text{cand}} \setminus \mathbf{x}$ ; // The set of  $N_q$  querying actions
- 9: // Evaluate these reference actions explicitly
- 10: **for each**  $x_i \in \mathbf{x}$  **do**
- 11:  $m_{\text{vir}} \leftarrow \text{Raycasting}(x_i, m_k)$ ;
- 12:  $I_i \leftarrow \text{ComputeCRMI}(m_{\text{vir}}, x_i)$ ;
- 13:  $\mathbf{y} \leftarrow \mathbf{y} \cup I_i$ ;
- 14: **end for**
- 15: // Find the suggested action using Algorithm 2
- 16:  $\{x_{\text{best}}, I_{\text{best}}\} \leftarrow \text{BKIOptimization}(\{\mathbf{x}, \mathbf{y}\}, \mathbf{x}^*, m_k)$ ;
- 17: **if**  $\text{Max}(I_{\text{best}}) > I_{\text{th}}$  **then**
- 18:  $x_{k+1} \leftarrow x_{\text{best}}(\text{MaxInfoIndex})$ ;
- 19:  $x_{\text{hist}} \leftarrow x_{\text{hist}} \cup x_{k+1}$ ;
- 20: **else**
- 21:  $x_{k+1} \leftarrow x_{k-1}$ ; // Back to previous action
- 22: Remove  $x_{k-1}$  from  $x_{\text{hist}}$ ;
- 23: **end if**
- 24: // Execute the action and update the map
- 25:  $\mathcal{P}_{\text{local}} \leftarrow \text{Astar}(x_k, x_{k+1})$  // Plan local path by A\*
- 26:  $m_{k+1} \leftarrow \text{ConfidenceRichMapping}(\mathcal{P}_{\text{local}})$ ;
- 27: **end while**

---

*Proposition 2:* The time complexity of our proposed BKIO method for evaluating all generated actions in Algorithm 1 and 2 is  $\mathcal{O}(N(N_z N_c^2 + N_z N_c^2 \lambda_z^{-1} \lambda_m^{-1}))$  of explicit CRMI evaluation and  $\mathcal{O}(N_{\text{epoch}} N \log N_q)$  of BKIO. The GPBO-based CRMI exploration methods in Section II-B have the same time complexity of explicit CRMI evaluation as BKIO but a complexity of  $\mathcal{O}(N_{\text{epoch}}(N^3 + N^2 N_q))$  to perform the expensive GP inference for CRMI, where  $N_{\text{epoch}}$  is the number of optimization epochs.

Comparatively, the NG-based exploration method which evaluates all actions explicitly has a time complexity of  $\mathcal{O}((N + N_q)(N_z N_c^2 + N_z N_c^2 \lambda_z^{-1} \lambda_m^{-1}))$  for explicit CRMI evaluation according to Proposition 1. The above comparative results indicate the BKIO-based exploration method outperforms GPBO and NG methods significantly in time efficiency, especially in large-scale and cluttered places that need more samples to evaluate rapidly (much greater  $N$  and  $N_q$ ).

#### IV. RESULTS AND DISCUSSIONS

In this section, we aim to validate the performance of the proposed BKIO method, compared with other famous exploration methods such as naive greedy-based (“NG”) [6], and GPBO with multiple epochs (“GPBO”), where GPBO replaces Line 16 of Algorithm 1 with the function ‘GPBOptimization’ using (3). NG evaluates all actions ( $\mathbf{x}$  and  $\mathbf{x}^*$ ) explicitly using Algorithm 1 (except Lines 15–16). Note that NG in [6] can be treated as the

---

**Algorithm 2: BKIOptimization( ).**


---

**Require:** Occupancy map at  $k$ th time step  $m_k$ , the set of explicitly evaluated samples  $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ , the set of querying actions  $\mathbf{x}^*$ , the number of optimization epoch  $N_{\text{epoch}}$ , factor  $\alpha$

- 1:  $x_{\text{best}} \leftarrow \{\}, f_{\text{best}}^I \leftarrow \{\}$ ;
- 2: **for each epoch do**
- 3: // Compute the kernel function using (4)
- 4:  $k \leftarrow \text{KernelFunction}(\mathbf{x}^*, \mathbf{x})$ ;
- 5: // Infer CRMI and uncertainty using (11)
- 6:  $\bar{k} \leftarrow \Sigma k, \bar{y} \leftarrow k \cdot \mathbf{y}$ ;
- 7:  $\mathbf{y}^* \leftarrow \bar{y}/\bar{k}, \text{cov}(\mathbf{y}^*) \leftarrow \Sigma/\bar{k}$ ;
- 8:  $f^I \leftarrow \mathbf{y}^* + \alpha^{1/2} \sigma_{\mathbf{y}^*}$ ;
- 9:  $[x_s, \text{idx}] \leftarrow \text{Max}(\{f^I, \mathbf{y}\})$ ;
- 10: **if**  $x_s \notin \mathbf{x}$  **then**
- 11: // Evaluate the CRMI of  $x_s$  explicitly
- 12:  $m_{\text{vir}} \leftarrow \text{Raycasting}(x_s, m_k)$ ;
- 13:  $y_s = \text{ComputeCRMI}(m_{\text{vir}}, x_s)$ ;
- 14: // Add into  $\mathcal{D}$  to improve prediction accuracy
- 15:  $\mathbf{x} \leftarrow \mathbf{x} \cup x_s, \mathbf{y} \leftarrow \mathbf{y} \cup y_s, \mathbf{x}^* \leftarrow \mathbf{x}^* \setminus x_s$ ;
- 16: **else**
- 17:  $y_s \leftarrow \mathbf{y}(\text{idx})$ ;
- 18: **end if**
- 19: // Save actions with highest objective function values
- 20:  $x_{\text{best}} \leftarrow x_{\text{best}} \cup x_s, I_{\text{best}} \leftarrow I_{\text{best}} \cup y_s$
- 21: **end for**
- 22: **return**  $x_{\text{best}}, I_{\text{best}}$

---

ground truth for verifying the CRMI prediction accuracy and exploration performance.

#### A. Synthetic Environments and Dataset Results

We run 20 Monte Carlo (MC) trials for all numerical simulations and dataset experiments on a laptop PC with a 2.6 GHz Intel i5-4210 M CPU and 8 G RAM. The robot poses are obtained by Gmapping [33]. Note that the man-made binary ground truth maps are used only to simulate the true world and generate virtual sensor observations for robot localization and mapping, and the robot cannot access these maps for exploration purposes, i.e., it still faces *unknown environments*.

The minimum information threshold is  $I_{\text{th}} = 2$  bits, and the trade-off factor is  $\alpha = 1$ . The  $\lambda_z$  and  $\lambda_m$  for CRMI calculation are both 0.1. We also choose the parameters of  $\zeta = 0.001$  and  $\sigma_\theta = 100$  for the BKIO method. We use 30 points (each point with eight directions uniformly) to generate candidate actions for each method, i.e.,  $N + N_q = 240, N_q = 2N$ . We set  $N_{\text{epoch}} = 30$  for GPBO and BKIO methods. Note that this value is chosen particularly and validated in an experimental study. More details are in [27]. We set an extreme case of numerous candidate actions to test the time-saving performance of the proposed methods.

To compare the exploration performance using different methods more quantitatively, referring to [4], [5], we prefer to use map entropy and coverage rate of MC results. In the map entropy and coverage rate figures (c.f. Figs. 3 and 6), the solid lines depict the means of MC trials of each method, and the shaded regions represent the standard deviations. Here we omit the standard deviation of NG for the figure’s readability.

*Results from structured/unstructured cases:* To simulate indoor and field scenes, we generate two synthetic maps with




Fig. 2. Examples of CRMI-based robot exploration in unknown synthetic environments. *Yellow square*: the start and end points. The red line starting with a red square represents a planned robot action/pose in SE(2), and *green square* is the local path between two actions.

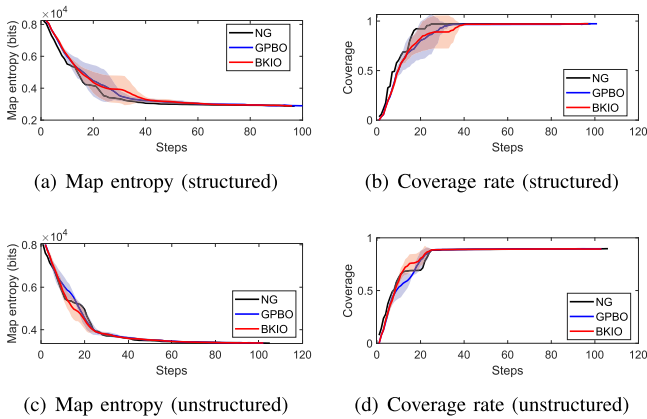


Fig. 3. Map entropy and coverage results of synthetic structured and unstructured maps. BKIO method can achieve similar exploration performance as GPBO, even better than GPBO in the unstructured map, and with less total exploration time.

the size of  $24 \text{ m} \times 14 \text{ m}$ , one structured maze map with several walls (Fig. 2(a),  $N_{\text{loop}} = 100$ ), and one unstructured map consisting of circles and ellipses representing forest-like scene (Fig. 2(b),  $N_{\text{loop}} = 150$ ). The unstructured one will obstruct the robot's FOV and need more exploration steps. The map resolutions are both  $0.2 \text{ m}$ . The simulated range sensor has a FOV of  $\pm 1.5 \text{ rd}$  with a resolution of  $0.05 \text{ rd}$  and a maximum sensing range of  $6 \text{ m}$ . The robot is initially at  $[1.2, 1.2] \text{ m}$  with  $0 \text{ rd}$  heading and trying to explore the prior unknown map.

The examples of resulting informative paths maximizing CRMI values are shown in Fig. 2. The quantitative results of structured and unstructured maps are shown in Fig. 3. BKIO and GPBO have similar evolution trends of map entropy reduction and coverage rate as NG in structured and unstructured maps, indicating the relatively low exploration performance loss in these BO-based methods.

*Results from dataset case:* To test our methods in a more realistic environment, we use a  $24 \text{ m} \times 14 \text{ m}$  Seattle map [34] containing narrow long corridors and cluttered rooms in Fig. 5(a). To implement the exploration methods, we apply the same simulated laser scanner as in the synthetic maps. The initial robot pose is  $x_0 = [11.4, 2.6, -\pi/2]$ . Fig. 5(a) selects one resulting informative path in Monte Carlo experiments. In the Seattle map, Fig. 6 shows BKIO and GPBO methods have similar reduction rates of map entropy and coverage rates w.r.t exploration steps. The simulation results provide evidence BKIO and GPBO can achieve comparable exploration performance in different environments.




Fig. 4. Exploration time cost using different methods in synthetic maps.




Fig. 5. BKIO-based robot exploration in the Seattle map. *Yellow square*: Start and end points. Note that the *y*-axis uses log10 representation.




Fig. 6. Map entropy and coverage results of the Seattle map. Note that the entropy and coverage would not be  $0 \text{ b}$  and  $100\%$ , respectively, since plenty of thick walls and obstacles are also counted, not only the inside free areas, as in Fig. 5(a).

*Time efficiency analysis:* As in Table I, we further analyze the time costs of both CRMI evaluation and exploration per iteration in each method. Both BKIO and GPBO are much more efficient than the NG method in each step, and both the time costs of CRMI evaluation and exploration get a significant reduction. In addition, BKIO shows higher time efficiency performance than GPBO in all scenes.

For example, compared to the NG method, BKIO decreases the information gain evaluation time by about  $90.4\%$  in the structured map, compared with the one of  $74.4\%$  using GPBO. These performance differences still prevail in unstructured and large cluttered maps, e.g. BKIO decreases the evaluation time costs per step by about  $92.6\%$  and  $93.9\%$  in the more complex maps of unstructured and Seattle, much better than  $76.1\%$  and  $77.4\%$  of GPBO, respectively. The box plots depicting the entire exploration time cost in Figs. 4 and 5(b) also support the above efficiency analysis.

*Local trade-off analysis:* An intuitive example is given to show how the UCB of (5) improves the local trade-off between exploration and exploitation at each exploration step. Due to the limited space, more details are presented in [27].

## B. Real-World Experiments

We deploy our proposed BKIO and GPBO on the Fetch<sup>2</sup> mobile robot running C++/Python code in the Robot Operating System (ROS). As shown in Fig. 1, we conduct experiments in an approximately  $25 \text{ m} \times 35 \text{ m}$  real-world office environment in Level 11, Building 2 of the University of Technology Sydney.

<sup>2</sup>[Online]. Available: <http://docs.fetchrobotics.com>

TABLE I  
 TIME COST COMPARISON PER STEP OF DIFFERENT METHODS

Methods	Structured map	Unstructured map	Seattle map [35]
NG [6]	28.31±1.49 / 51.05±2.66	40.59±3.04 / 71.72±3.27	53.02±5.52 / 101.15±3.11
GPBO	7.25±0.98 / 10.27±1.07	9.70±1.31/ 17.08±2.24	11.97±1.10 / 21.65±1.94
<b>BKIO</b>	<b>2.71±0.36 / 5.14±0.48</b>	<b>2.98±0.43/ 9.95±1.17</b>	<b>3.22±0.42 / 10.54±1.09</b>

Note: Time cost (in seconds) of CRMI evaluation per step / Time cost of exploration per step.

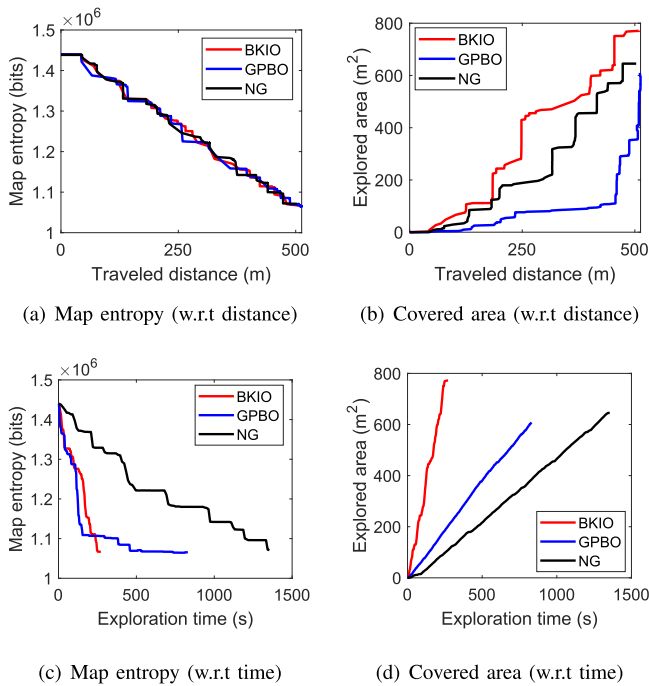


Fig. 7. Real-world map entropy and coverage results w.r.t robot travel distance and total time cost, respectively.

There are many chairs, tables, and glass walls in the office and kitchen zones, which would hinder the ego perception and motion of the robot. The FOV of the LiDAR is 220 degrees, and the max sensing range is 25 m. Note that the robot explores the environment autonomously and only uses the laser scan data from the SICK 2D LiDAR (15 Hz) without cameras. The max sensing range is set to 8 m, forcing this ‘sensing-limited’ robot to explore longer.

Here we use the Cartographer [35] as the localization module and CRM [8] as the mapping module in the exploration framework. The robot calls for a self-defined ROS service to evaluate the CRMI of each action generated at the extracted frontiers, then decides where to go at the next step until the exploration ends when the travel cost reaches 500 m or the expected CRMI of the best action is lower than the predefined threshold of 10 bits. An example video is available here.<sup>3</sup>

Fig. 7 shows the comparative evolution results of map entropy and explored areas using three different methods as the travel distance and total exploration time increase. The map entropy and coverage curves in Fig. 7(a) and (b) show these methods all complete the exploration finally. In Fig 7(a), the map entropy curves of the three methods w.r.t the traveled distance are similar, but the explored areas vary a lot in Fig. 7(b). This is mainly because we determine a map cell as already explored if its expected occupancy deviates from 0.5, otherwise unexplored (bool 0 or

1) when computing the covered area; but when computing the map entropy, we use the continuous expected occupancy (in the interval [0, 1]). The rapid increase of explored areas in Fig. 7(b) and (d) also implies our BKIO can drive the robot to explore more unknown areas rapidly, given a limited distance budget, and get an environmental sketch for the subsequent robot application.

Fig. 7(b) shows that the GPBO travels a longer distance at the start stage but gets the less explored area than others since it moves in the corridors for a long distance and explores more aggressively and roughly in the later stage. In Fig. 7(d), we can observe that the BKIO and GPBO methods outperform the NG in terms of exploration time cost significantly (79.9% and 38.6% reduction, respectively), where BKIO performs better in explored areas than the others. Fig. 7(c) and (d) also show our methods conduct much more efficient exploration in map uncertainty reduction and exploring unknown areas.

Moreover, as in Fig. 7(d), BKIO reduces about 67.3% total time cost to end the exploration and increases about 21% covered areas than GPBO. Further, the mean time costs of each exploration step are 12.2095±0.6383 s, 0.2007±0.1362 s, and 0.0038±0.0033 s for NG, GPBO, and BKIO, respectively. This also evidences the advantage of our proposed methods.

The overall exploration sequences of the three methods also show some differences, especially at the intersections and cluttered areas. This mainly attributes to the unequal prediction accuracy at each step and the difference accumulation over time. More analyses are in [27].

### C. Discussion

As presented in previous experiments of typical scenes, our proposed methods (GPBO and BKIO) have shown desired exploration performance in efficiency and accuracy, especially the more efficient BKIO. Moreover, we use the public packages (e.g., scikit-learn and NumPy) to implement GP, but the BKI code is non-optimized. Hence we believe our methods’ above performance improvement trend is not biased and would be more salient when facing more complex scenes.

The significant time-saving benefits for online exploration tasks greatly, especially when using computationally limited or energy-limited platforms. In other words, we offer two options for efficient information-based robot exploration tasks in unstructured and relatively large areas.

Another thing we should clarify is that we are not aiming to propose a new exploration framework. The exploration framework we use in this letter (i.e., Algorithm 1) stems from depth-first search; it is just used to run BKIO, GPBO, and NG methods. In other words, we can use another exploration framework instead, such as the sampling-based incrementally-informative graph [10]. We have also tested our methods in both sampling-based (Section IV-A) and frontier-based (Section IV-B) exploration. Hence we would not discuss the advantages and disadvantages of these frameworks because we focus on decreasing the evaluation time cost at each exploration step and the time cost of the whole exploration.

<sup>3</sup>[Online]. Available: <https://youtu.be/sOW4fuaAwT8>

Moreover, the information metric is also not limited to CRMI; one can apply our methods to other information metrics for robot exploration, such as CSQMI [11] and GP-based MI (GPMI) [10].

In some experiments (e.g., Fig. 3), one can also find NG sometimes performs worse during the exploration (14–21 steps). The underlying reason is mainly that NG completely pursues the action maximizing the expected CRMI (i.e., *exploitation*), but may neglect more promising actions that have less expected MI currently (i.e., *exploration*). Our future work will also study the non-myopic informative path planning to realize the global trade-off between exploration and exploitation.

## V. CONCLUSION

This letter mainly contributed to new learning-based approaches with high efficiency for information-theoretic robot exploration in unknown environments. In particular, an information gain evaluation method for inferring the CRMI of numerous sampled robot actions is proposed based on GP and BO. With the Bayesian kernel inference-based BO, the time complexity of CRMI prediction decreases to a logarithm level. An informative, objective function integrating the predicted CRMI and uncertainty is also used to ensure the bounded regret and local trade-off between exploration and exploitation at each step. The proposed method also gets verified under an autonomous exploration framework by extensive simulations and real-world experiments in different scenes. This reveals that the proposed BKIO and GPBO exploration methods outperform the non-learning ones overall in efficiency without losing much exploration performance, especially in unstructured and large cluttered scenes. This would also benefit other information-based exploration methods stuck by efficiency issues. Future work mainly involves the extension to 3D exploration.

## REFERENCES

- [1] H. Azpúrua, M. F. M. Campos, and D. G. Macharet, “Three-dimensional terrain aware autonomous exploration for subterranean and confined spaces,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2021, pp. 2443–2449.
- [2] P. Stankiewicz, Y. T. Tan, and M. Kobilarov, “Adaptive sampling with an autonomous underwater vehicle in static marine environments,” *J. Field Robot.*, vol. 38, no. 4, pp. 572–597, 2021.
- [3] B. J. Julian, S. Karaman, and D. Rus, “On mutual information-based control of range sensing robots for mapping applications,” *Int. J. Robot. Res.*, vol. 33, no. 10, pp. 1375–1392, 2014.
- [4] B. Charrow, S. Liu, V. Kumar, and N. Michael, “Information-theoretic mapping using Cauchy-Schwarz quadratic mutual information,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2015, pp. 4791–4798.
- [5] Z. Zhang, T. Henderson, S. Karaman, and V. Sze, “FSMI: Fast computation of Shannon mutual information for information-theoretic mapping,” *Int. J. Robot. Res.*, vol. 39, no. 9, pp. 1155–1177, 2020.
- [6] Y. Xu, R. Zheng, M. Liu, and S. Zhang, “CRMI: Confidence-rich mutual information for information-theoretic mapping,” *IEEE Robot. Automat. Lett.*, vol. 6, no. 4, pp. 6434–6441, Oct. 2021.
- [7] Y. Xu, R. Zheng, S. Zhang, and M. Liu, “Confidence-rich localization and mapping based on particle filter for robotic exploration,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 4471–4477.
- [8] A.-A. Agha-Mohammadi, E. Heiden, K. Hausman, and G. Sukhatme, “Confidence-rich grid mapping,” *Int. J. Robot. Res.*, vol. 38, no. 12/13, pp. 1352–1374, 2019.
- [9] G. A. Hollinger and G. S. Sukhatme, “Sampling-based robotic information gathering algorithms,” *Int. J. Robot. Res.*, vol. 33, no. 9, pp. 1271–1287, 2014.
- [10] M. G. Jadidi, J. V. Miró, and G. Dissanayake, “Sampling-based incremental information gathering with applications to robotic exploration and environmental monitoring,” *Int. J. Robot. Res.*, vol. 38, no. 6, pp. 658–685, 2019.
- [11] B. Charrow, V. Kumar, and N. Michael, “Approximate representations for multi-robot control policies that maximize mutual information,” *Auton. Robot.*, vol. 37, no. 4, pp. 383–400, 2014.
- [12] B. Charrow et al., “Information-theoretic planning with trajectory optimization for dense 3D mapping,” in *Proc. Robot. Sci. Syst.*, 2015, vol. 11, pp. 3–12.
- [13] S. Bai, F. Chen, and B. Englot, “Toward autonomous mapping and exploration for mobile robots through deep supervised learning,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2017, pp. 2379–2384.
- [14] F. Chen, J. D. Martin, Y. Huang, J. Wang, and B. Englot, “Autonomous exploration under uncertainty via deep reinforcement learning on graphs,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2020, pp. 6140–6147.
- [15] T. Wang, R. Liao, J. Ba, and S. Fidler, “NerveNet: Learning structured policy with graph neural networks,” in *Proc. Int. Conf. Learn. Representation*, 2018.
- [16] R. Marchant and F. Ramos, “Bayesian optimisation for informative continuous path planning,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2014, pp. 6136–6143.
- [17] R. Oliveira, L. Ott, V. Guizilini, and F. Ramos, “Bayesian optimisation for safe navigation under localisation uncertainty,” in *Proc. Int. Symp. Robot. Res.*, 2020, pp. 489–504.
- [18] G. Francis, L. Ott, R. Marchant, and F. Ramos, “Occupancy map building through Bayesian exploration,” *Int. J. Robot. Res.*, vol. 38, no. 7, pp. 769–792, 2019.
- [19] S. Bai, J. Wang, K. Doherty, and B. Englot, “Inference-enabled information-theoretic exploration of continuous action spaces,” in *Proc. Int. Symp. Robot. Res.*, 2015, pp. 419–433.
- [20] S. Bai, J. Wang, F. Chen, and B. Englot, “Information-theoretic exploration with Bayesian optimization,” in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2016, pp. 1816–1822.
- [21] W. R. Vega-Brown, M. Doniec, and N. G. Roy, “Nonparametric Bayesian inference on multivariate exponential families,” in *Proc. Int. Conf. Neural Inf. Process. Syst.*, vol. 27, 2014.
- [22] V. Peretroukhin, W. Vega-Brown, N. Roy, and J. Kelly, “PROBE-GK: Predictive robust estimation using generalized kernels,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 817–824.
- [23] C. Richter, W. Vega-Brown, and N. Roy, “Bayesian learning for safe high-speed navigation in unknown environments,” in *Proc. Int. Symp. Robot. Res.*, 2015, pp. 325–341.
- [24] T. Shan, J. Wang, B. Englot, and K. Doherty, “Bayesian generalized kernel inference for terrain traversability mapping,” in *Proc. Conf. Robot Learn.*, 2018, pp. 829–838.
- [25] K. Doherty, T. Shan, J. Wang, and B. Englot, “Learning-aided 3-D occupancy mapping with Bayesian generalized kernel inference,” *IEEE Trans. Robot.*, vol. 35, no. 4, pp. 953–966, Aug. 2019.
- [26] L. Gan, R. Zhang, J. W. Grizzle, R. M. Eustice, and M. Ghaffari, “Bayesian spatial kernel smoothing for scalable dense semantic mapping,” *IEEE Robot. Automat. Lett.*, vol. 5, no. 2, pp. 790–797, Apr. 2020.
- [27] Y. Xu, R. Zheng, S. Zhang, M. Liu, and S. Huang, “CARE: Confidence-rich autonomous robot exploration using Bayesian kernel inference and optimization,” 2023. [Online]. Available: [https://github.com/Shepherd-Gregory/BKIO-Exploration/blob/main/RAL2023\\_full\\_version.pdf](https://github.com/Shepherd-Gregory/BKIO-Exploration/blob/main/RAL2023_full_version.pdf)
- [28] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. Cambridge, MA, USA: MIT Press, 2005.
- [29] S. T. O’Callaghan and F. T. Ramos, “Gaussian process occupancy maps,” *Int. J. Robot. Res.*, vol. 31, no. 1, pp. 42–62, 2012.
- [30] N. Srinivas, A. Krause, S. Kakade, and M. Seeger, “Gaussian process optimization in the bandit setting: No regret and experimental design,” in *Proc. Int. Conf. Mach. Learn.*, 2010, pp. 1015–1022.
- [31] P. Auer, N. Cesa-Bianchi, and P. Fischer, “Finite-time analysis of the multiarmed bandit problem,” *Mach. Learn.*, vol. 47, no. 2, pp. 235–256, 2002.
- [32] L. Gan et al., “Multitask learning for scalable and dense multilayer Bayesian map inference,” *IEEE Trans. Robot.*, vol. 39, no. 1, pp. 699–717, Feb. 2023.
- [33] G. Grisetti, C. Stachniss, and W. Burgard, “Improved techniques for grid mapping with Rao-Blackwellized particle filters,” *IEEE Trans. Robot.*, vol. 23, no. 1, pp. 34–46, Feb. 2007.
- [34] A. Howard and N. Roy, “The robotics data set repository (Radish),” 2003. [Online]. Available: <http://radish.sourceforge.net/>
- [35] W. Hess, D. Kohler, H. Rapp, and D. Andor, “Real-time loop closure in 2D LIDAR SLAM,” in *Proc. IEEE Int. Conf. Robot. Automat.*, 2016, pp. 1271–1278.